



# Goodfact Security Policies Whitepaper

Last Updated: September 2024

|                                                              |          |
|--------------------------------------------------------------|----------|
| <b>1. Data Residency, Privacy and Asset Management .....</b> | <b>2</b> |
| Data Location .....                                          | 2        |
| Data Transportation.....                                     | 2        |
| Physical Environment.....                                    | 2        |
| Data Delivery and Access .....                               | 2        |
| Description of Supporting Service Architecture .....         | 3        |
| Third-Party Integrations.....                                | 3        |
| <b>2. Organization Governance .....</b>                      | <b>4</b> |
| Confidentiality and Encryption.....                          | 4        |
| Data Handling.....                                           | 4        |
| Change Management Process.....                               | 4        |
| Technical Delivery Process .....                             | 4        |
| Automated Testing .....                                      | 5        |
| Incident Report Process .....                                | 5        |
| <b>3. Security .....</b>                                     | <b>6</b> |
| System Procedures and Processes .....                        | 6        |
| Organizational and Technical Security .....                  | 6        |
| Security Patches and Upgrades .....                          | 6        |
| <b>4. Data Storage.....</b>                                  | <b>6</b> |
| Loss Prevention.....                                         | 6        |
| Data Expungement.....                                        | 7        |
| Capacity and Performance.....                                | 7        |
| Baseline Configuration .....                                 | 7        |
| Vulnerability/Patch Management.....                          | 7        |
| Disaster Recovery .....                                      | 7        |
| <b>5. Service Level Objectives.....</b>                      | <b>8</b> |
| Uptime and Platform Availability.....                        | 8        |
| Customer Support Scope and Availability Guidelines .....     | 8        |



## 6. Generative AI Usage .....8

# 1. Data Residency, Privacy and Asset Management

## Data Location

Goodfact utilizes servers that are exclusively positioned within Canada, leveraging the Central Canada Region of Amazon Web Services (AWS). This ensures our adherence to Canadian privacy laws and regulations in delivering our services.

## Data Transportation

To ensure secure communication, we employ the highest level of Transport Layer Security (TLS, commonly referred to as SSL) encryption. This covers communication between your browser and our application, as well as internal communication between our systems, providing end-to-end encryption. Our services are hosted in a Virtual Private Network (VPC) that blocks any communication to and from the internet.

## Physical Environment

Our application containers run within a secure physical environment in Amazon's data centers. Amazon data centers are best-in-class in physical security, employing intrusion detection systems, CCTV, security staff and multi-factor authentication, among many other security measures. You can read more about them here: <https://aws.amazon.com/compliance/data-center/controls/>

Our application containers are not directly accessible from the internet and are concealed behind Application Load Balancers. Instead of running servers, we use Amazon Elastic Container Service (ECS) for application containers and Amazon Relational Database Service (RDS) for data hosting. Our files are stored in Amazon Simple Storage Solution (S3)

## Data Delivery and Access

Operating in a secure environment, Goodfact segregates production from testing and development environments. This segregation is facilitated by AWS at the highest level available (separate accounts). Only the CEO and CTO are granted access to the production environment, with all actions meticulously logged and monitored for enhanced security.

We strictly deliver data to authorized web clients via SSL/TLS (also known as HTTPS). Any other form of communication to our system is not possible. We utilize the most robust TLS



standard currently supported by the majority of web browsers (TLS1.3 and TLS1.2 with only strong cipher suites enabled).

## Description of Supporting Service Architecture

Our supporting service architecture is composed of encrypted databases. All data, including files and database backups, are encrypted at-rest using AES256, one of the most robust and well-tested encryption standards. To thwart man-in-the-middle attacks, communication between clients and our services is encrypted using HTTPS. HTTPS is terminated by Application Load Balancers, and all services operate in Virtual Private Clouds that are isolated from the internet (no access to or from the internet).

We perform updates to our API components by constructing a new containerized image and pushing it to AWS Elastic Container Registry. The AWS Elastic Container Service then loads the new image and starts using it in a blue/green deploy.

## Third-Party Integrations

Beyond Amazon Web Services, Goodfact does not employ any third-party services. All monitoring and analytics are executed by self-hosted services in our secure VPC or by AWS services.

For secure payment processing, we utilize Stripe. However, it operates strictly in-browser, and Stripe has no access to any backend resources. Stripe is a highly regarded payment processor and is PCI and SOC-2 compliant. You can read more about Stripe's security measures in their security documentation: <https://stripe.com/docs/security>



## 2. Organization Governance

### Confidentiality and Encryption

Goodfact adheres to rigorous confidentiality protocols. Access to the production system, which stores client data, is strictly limited to the CEO and CTO. The encryption keys used for data security are stored in an ISO27001 certified facility (AWS Key Management Service and AWS Secrets Manager) and are never accessed by employees, consultants, or third parties.

### Data Handling

The production environment, which houses sensitive data, is strictly accessible only by the CEO and CTO of Goodfact. No other employees, consultants, or third parties have access to this environment.

The production environment operates in a separate AWS account, which significantly minimizes the chances of accidental or intentional data leakage. All modifications to the production environment are performed via automated and access-controlled systems.

### Change Management Process

The change management process is meticulously monitored with a strong emphasis on security considerations. All changes undergo a code-review process on Github, which tracks the changes made, by whom, and when. Github also preserves all previous versions of the code, enabling us to analyze changes and revert them easily if necessary.

### Technical Delivery Process

Following code review, a new containerized image is constructed using Docker. The image is uploaded to AWS Elastic Container Registry, and our Elastic Container Service (ECS) is updated to switch to the new image. AWS ECS performs a blue/green deploy - it starts the new image, performs health checks, and once they succeed, starts routing traffic to the new image.

Since AWS Elastic Container Registry retains all previous versions of our containerized images, we can effortlessly switch back to a previous version should any issues arise.

Front-end releases are similarly performed by constructing new front-end assets using Node.js, which are then uploaded to AWS S3. We use AWS CloudFront to serve our front-end assets. All front-end releases are namespaced, enabling us to easily switch to different versions of our front-end application if any issues arise with the release.

Our technical infrastructure is managed using AWS CloudFormation, and no manual changes to our AWS infrastructure are permitted. All changes must be executed via CloudFormation.



## Automated Testing

We employ automated unit testing and integration testing to verify new back-end releases. We have comprehensive end-to-end tests verifying our front-end releases and ensuring the front-end is communicating with the back-end without issues.

## Incident Report Process

In the event of a security incident, Goodfact's CEO and CTO are immediately mobilized, initiating a comprehensive investigation. The priority is to secure data and contain potential data leaks, even if necessitating a temporary application shutdown.

Once containment measures are activated, customers impacted by the incident are promptly notified. The investigative team then determines the specifics of the breach, including the nature of the compromised data, the affected parties, and the duration and extent of the data leak. Regular updates are provided to customers throughout this process.

Subsequently, the team identifies the root causes of the incident. A strategy is then formulated to restore system functionality while addressing the root causes. If the incident necessitates long-term changes, a secure method to restore the system is devised, potentially with limited functionality initially. A roadmap for these changes is developed and communicated to the affected customers.

Upon resolution, a detailed report is prepared, outlining the incident's nature, scope, the affected parties, the remediation steps taken, the reasons for the system's current secure status, and the future preventive measures. This report is shared with customers, reinforcing our commitment to transparency, data security, and customer trust at Goodfact.



## 3. Security

### System Procedures and Processes

At Goodfact, we adopt a security-in-layers and minimum-access approach. Our systems are decoupled and have exclusive access to only the resources they need to function. This reduces the risk of accidentally exposing data. The production environment is only accessible by the CEO and CTO, which reduces the risk of unauthorized access. Furthermore, our use of AWS tooling and self-hosted functionality guarantees that no data leaves our AWS account.

### Organizational and Technical Security

Access control at Goodfact is managed using AWS IAM with 2-factor-authentication using TOTP to prevent SMS-based attacks. Access to the production account is restricted to only the CEO and CTO. No employees, consultants, or third parties have access to the production account. Despite being a small team, we ensure all personnel are thoroughly versed in security best practices.

Our architecture has been reviewed by an AWS Senior Solutions Architect with a focus on security. We utilize AWS GuardDuty for continuous monitoring of any security concerns. We use password authentication sparingly, instead opting for key-based authentication where possible, or using build systems to execute changes instead of manually logging into systems. Where we do have to store passwords we use KeePassXC to securely store them. All access to any system - production or development - is logged and audited regularly.

### Security Patches and Upgrades

As part of our commitment to security, we regularly apply patches and upgrades to keep our systems up-to-date with the latest security measures. Since we don't operate any servers our update process is streamlined - a new image is built and released. We apply security updates as soon as we receive notification, and we update our technical dependencies regularly (both for front-end and back-end builds). We use Github Dependabot to monitor any security concerns that would apply to our dependencies.

## 4. Data Storage

### Loss Prevention

To prevent data loss, Goodfact operates on an append-only data policy where no data is deleted. We only mark data as deleted but never actually delete it, unless explicitly requested by a customer on account closing. This minimizes the risk of accidental data loss.



## Data Expungement

In case of a specific request from a customer on account closure, we have a meticulous process in place to expunge their data from our systems. We will ensure that all customer data is removed from our database and our file storage. We will also ensure that customer data is removed from all backups within 90 days.

## Capacity and Performance

Our servers are designed to handle high capacity loads without compromising performance or security. The systems are built on AWS infrastructure, ensuring high reliability and availability. With AWS Elastic Container Service we can easily add additional resources, and the AWS RDS system can be scaled vertically to address any customer needs. We track performance by measuring how long it takes our front-end to load and display data. Our goal is to have a chronology load within 2 seconds.

## Baseline Configuration

Our baseline configuration includes stringent firewall rules, secure application load balancers, and strong encryption standards.

## Vulnerability/Patch Management

We regularly self-audit and monitor access logs. Furthermore, we use AWS GuardDuty to detect any suspicious activity on our systems. Regular patches and upgrades are applied to maintain a robust security posture. Our container-based architecture makes it easy to quickly apply security updates without manual intervention. We never manually access the production containers.

## Disaster Recovery

In case of a disaster, our recovery plan involves spinning up another instance of Goodfact in a different region using CloudFormation. Our use of highly-distributed (S3) and highly-reliable managed database (AWS RDS) systems ensures resilience against regional disasters by storing data across multiple availability zones.



## 5. Service Level Objectives

### Uptime and Platform Availability

Goodfact is committed to providing high uptime and platform availability, ensuring that your operations can run smoothly without interruption. We target a 99.9% uptime by utilizing a highly-resilient containerized architecture that removes deployment downtime with blue/green deployment approach. And should any issues occur after a release, we can revert to a previous release within minutes.

### Customer Support Scope and Availability Guidelines

Our customer support team is always ready to assist you with any security-related concerns or issues you may encounter while using our platform. You can reach us via email at [info@goodfact.co](mailto:info@goodfact.co). For technical questions, you can contact our CTO Thomas Rubbert directly at [thomas@goodfact.co](mailto:thomas@goodfact.co). For general questions, you can contact our CEO Tali Green at [tq@goodfact.co](mailto:tq@goodfact.co).

## 6. Generative AI

Goodfact employs a **closed-loop system** for handling user data. This means that the data we receive from our users is processed solely for the intended tasks and never used to train or fine-tune any underlying machine learning models, including the large language models (LLMs) that we use through AWS Bedrock.

Unlike **open-loop systems**—where user data can be integrated back into the model for further training or optimization—our closed-loop approach ensures that user data remains fully isolated from model development pipelines. This guarantees that your data is processed only within the scope of the services we provide, protecting it from being repurposed or shared.